

WE CLAIM:

1. A computer system for generating metadata for use during stack unwinding, comprising:
  - a plurality of procedures wherein each procedure comprises a sequence of binary instructions;
  - a first plurality of blocks of metadata wherein each block of metadata is associated with a corresponding procedure in the plurality of procedures; and
  - an unwind rewriter programmed to generate a second plurality of blocks of metadata from the first plurality of blocks of metadata in response to a modification of the sequence of binary instructions within a procedure, such that the second plurality of blocks of metadata accurately represents the modified sequence of binary instructions.
2. The computer system of claim 1, wherein each block of metadata in the plurality of blocks of metadata includes at least one unwind table and at least one unwind information block.
3. The computer system of claim 2, wherein the at least one unwind information block includes a region header describing a region of zero length.
4. A computer-implemented method of regenerating unwind data for a modified binary procedure wherein a current order of basic blocks within the modified binary procedure differs from an original order of the basic blocks, the computer-implemented method comprising:
  - discovering original unwind data that describes the original order of the basic blocks;
  - regenerating new unwind data, based in part on information in the original unwind data, that represents the current order of basic blocks within the modified binary procedure; and
  - writing the new unwind data to the modified binary procedure.

5. The method of claim 4, wherein discovering the unwind data comprises parsing the original unwind data that describes the original order of the basic blocks.

6. The method of claim 5, wherein parsing the original unwind data comprises identifying a start basic block and an end basic block of a region associated with the modified binary procedure.

7. The method of claim 6, wherein identifying the end basic block of the region further comprises splitting a single basic block into two basic blocks, such that a first basic block ends on a last instruction of the region.

8. The method of claim 6, wherein parsing the original unwind data further comprises identifying an unwind information block associated with a basic block in the original order of the basic blocks that includes a when action description record and establishing a link between the when action description record and the corresponding instruction in the basic block.

9. The method of claim 4, wherein regenerating new unwind data comprises regenerating new unwind tables and new unwind descriptor records.

10. The method of claim 9, wherein regenerating the new unwind descriptor records further comprises determining if basic blocks identified in a single unwind table associated with the original order of basic blocks are associated with more than one unwind table associated with the current order of basic blocks, and if so, creating a new region header describing a region of zero length.

11. A computer-implemented method for regenerating unwind data in response to a binary modification to a procedure, the procedure including a plurality of basic blocks, comprising:

receiving a first group of unwind data comprising an unwind table and a plurality of unwind descriptor records wherein the first group of unwind data is associated with the unmodified procedure;

parsing the first group of unwind data to identify a start basic block and an end basic block for a region associated with the procedure; and

generating a second group of unwind data comprising a second unwind table and a second plurality of unwind descriptor records such that the second group of unwind data accurately represents the binary modification to the procedure.

12. The computer-implemented method of claim 11, wherein parsing the first group of unwind data further comprises generating a relationship between a when-descriptor within an unwind descriptor record and an instruction in the procedure, and generating the second group of unwind data further comprises associating the when-descriptor with an appropriate unwind descriptor record in the second plurality of unwind descriptor records.

13. The method of claim 11, wherein identifying the end basic block of the region further comprises splitting a single basic block into two basic blocks, such that a first basic block ends on a last instruction of the region.

14. The method of claim 11, wherein generating the second plurality of unwind descriptor records further comprises determining if basic blocks identified in a single unwind table associated with the unmodified procedure are associated with more than one unwind table associated with the binary modified procedure, and if so, creating a new region header describing a region of zero length.

15. A computer-readable medium encoded with a data structure for providing metadata representing a binary program including basic blocks that have been modified from a first order to a second order, the data structure comprising:

a first metadata field associated with a first group of instructions within a first procedure; and

a second metadata field associated with a second group of instructions within the first procedure, wherein the first metadata field and the second metadata field accurately reflect a flow of instructions of the basic blocks in the second order.

16. The data structure of claim 15, further comprising a third group of instructions associated with a second procedure, wherein the third group of instructions lies within the flow of instructions between the first group of instructions and the second group of instructions.

17. The data structure of claim 15, wherein each metadata field includes at least one unwind table and at least one unwind information block.

18. The computer system of claim 15, wherein the second metadata field includes a region header describing a region of zero length.

19. A computer-readable medium having computer-executable instructions for regenerating unwind data in response to a binary modification to a procedure, the procedure including a plurality of basic blocks, the instructions comprising:

receiving a first group of unwind data comprising an unwind table and a plurality of unwind descriptor records wherein the first group of unwind data is associated with the unmodified procedure;

parsing the first group of unwind data to identify a start basic block and an end basic block for a region associated with the procedure; and

generating a second group of unwind data comprising a second unwind table and a second plurality of unwind descriptor records such that the second group of unwind data accurately represents the binary modification to the procedure

20. A computer-readable medium having computer-executable instructions for generating unwind data for a modified binary procedure wherein a current order of basic blocks within the modified binary procedure differs from an original order of the basic blocks, the instructions comprising:

discovering original unwind data that describes the original order of the basic blocks;

regenerating new unwind data, based in part on information in the original unwind data, that represents the current order of basic blocks within the modified binary procedure; and

writing the new unwind data to the modified binary procedure.